H-1131

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Appl. No. 10/780,772 Confirmation No. 2065

Inventor: SHIMADA, K. et al

Filed: February 19, 2004

For: STORAGE HAVING LOGICAL PARTITIONING CAPABILITY AND SYSTEMS WHICH INCLUDE THE STORAGE

Group Art Unit: 2186

Examiner: Unassigned

Docket No.: H-1131

Customer No.: 24956

PETITION TO MAKE SPECIAL
UNDER 37 CFR §1.102(d) (MPEP §708.02(VIII))

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450 November 30, 2004

Sir:

The Applicants petition the Commissioner to make the above-identified application special in accordance with 37 CFR §1.102(d). In support of this Petition, pursuant to MPEP § 708.02(VIII), Applicants state the following.

(A) REQUIRED FEE

This Petition is accompanied by the fee set forth in 37 CFR § 1.117(h). A Credit Card Payment Form in the amount of

$130.00 accompanies this Petition in satisfaction of the fee. The Commissioner is hereby authorized to charge any additional payment due, or to credit any overpayment, to Deposit Account No. 50-1417.

(B) ALL CLAIMS DIRECTED TO A SINGLE INVENTION

All the pending claims of the application, claims 1-18, are directed to a single invention. If the Office determines that all claims in the application are not directed to a single invention, Applicant will make election without traverse as a prerequisite to the grant of special status.

The claimed invention is directed to a storage system, connected to a network, that includes a plurality of interfaces which are connected to the network, and which receive file access. The system also includes a plurality of disk drives and a control unit which translates the file access into block access and controls the plurality of disk drives on the basis of the block access. The control unit logically partitions the plurality of interfaces, the plurality of disk drives, and the control unit, and causes the partitioned plurality of interfaces, the partitioned plurality of disk drives, and the partitioned control unit to operate independently as a plurality of virtual storages.

Under additional aspects, the storage system control unit may include a plurality of cache memories, and the plurality of cache memories is logically partitioned and allocated to the respective plurality of virtual storages. The control unit may further include a first processor that translates the file access into the block access, and a second processor that controls the plurality of disk drives on the basis of the block access. The first processor and the second processor are logically partitioned and allocated to the respective plurality of virtual storages. A plurality of memories may also be included and used by the first processor, and a plurality of communication units may connect the first processor and the second processor, and the memories and communication units may also be logically partitioned by the control unit.

Under yet another aspect, the storage system includes a supervising terminal connected to the storage, and the storage logically partitions the plurality of interfaces, the plurality of disk drives, and the control unit on the basis of the information to be inputted to the supervising terminal, and operates as plural virtual storages independently. The information to be inputted to the supervising terminal may be information on characteristics of accesses of a computer using

the storage, and the storage may calculate an amount of logical partitioning of resources provided in the storage on the basis of the information on characteristics of accesses to be inputted to the supervising terminal, and performs the logical partitioning as a result of the calculation.

(C) PRE-EXMINATION SEARCH

A careful and thorough pre-examination search has been conducted, directed to the invention as claimed in claims 1-18 of the application. The pre-examination search was conducted in the following areas: Class 711, subclasses 147, 153, 173, and 202; and Class 718, subclass 1. Additionally, a computer database search was conducted on the USPTO EAST database.

(D) REFERENCES MOST CLOSELY RELATED TO THE CLAIMED INVENTION

Of the documents reviewed during the search, those deemed to be most closely related to the subject matter encompassed by the claims are listed below. These documents were made of record in the present application by the Information Disclosure Statement (IDS) filed November 24, 2004.

Patent No.	Inventor(s)
US 5,210,844	Shimura et al.
US 5,592,638	Onodera
US 5,659,786	George et al.
US 5,704,055	George et al.

US 5,790,852	Salm
US 6,484,245	Sanada
US 6,606,690	Padovano
US 6,640,278	Nolan et al.
US 6,725,352	Goodman et al.
US 6,738,854	Hoese et al.
US 6,742,090	Sanada et al.
US 6,754,776	Conway et al.
US 6,763,419	Hoese et al.

Pub. Pat. Appl.	Inventor(s)
US 20020007366	Fontijn
US 20020133539	Monday
US 20030065898	Flamma et al.
US 20030120751	Husain et al.
US 20040111580	Weber et al.

Additionally, the following documents were made of record in the present application by the IDS filed February 19, 2004.

Pub. Pat. Appl.	Inventor(s)
US20030097393	Kawamoto et al.
JP2003-157177	Kawamoto et al.

Because all of the above-listed documents have been made of record in the present application by the IDSs filed February 19, 2004, and November 24, 2004, additional copies of the documents have not been submitted with this Petition in accordance with MPEP § 708.02(VIII)(D).

(E) DETAILED DISCUSSION OF THE REFERENCES

A discussion of each the above-listed documents is set forth below, pointing out, with the particularity required by

37 CFR 1.111 (b) and (c), how the claimed subject matter is patentable over the teachings of the above-listed documents.

US 5,210,844, Shimura et al., discloses an information processing apparatus that has at least one processor and a main storage, accessed by the processor, and capable of providing a plurality of logical information processing apparatuses by logically partitioning the information processing apparatus. The information processing apparatus includes a main storage partitioned into a plurality of memory areas, with each of the memory areas corresponding to one of the plurality of logical information processing apparatuses. The information processing apparatus further includes a first storage unit for storing identification information for each of the memory areas identifying the logical information processing apparatus allocated to each memory, and a read unit for reading the identification information from the first storage unit when the main storage is to be accessed by one of the plurality of logical information processing apparatuses. (See, e.g., Column 3, line 27 - Column 4, line 40.) However, Shimura fails to teach the present invention, including a supervising terminal connected to the storage, so that the storage logically partitions a plurality of interfaces, a

plurality of disk drives, and a control unit on the basis of information to be inputted to the supervising terminal.

US 5,592,638, Onodera, discloses a data processor having a storage and a processor, and a method of assigning a plurality of regions of the storage to a plurality of virtual machines according to a plurality of activating or non-activating storage region assignment requests. A virtual machine control program (VMCP) 110 is operated on a real data processor 10. Under control of the VMCP 110, a plurality of virtual machines or logical partitions are generated and the respective logical partitions are assigned to the storage regions of a real storage 20. In particular, the logical partition whose storage origin is not designated is activated, the VMCP 110 reads a storage assignment status table 310 of the file memory 30, generates a storage assignment work table 410 and the remainder storage region area table 420 on a work memory 40, and determines whether or not assignment of the storage region of the associated logical partition is possible. (See, e.g., Column 4, lines 58-67.) Thus, while Onodera is concerned with assigning a plurality of logical partitions, Onodera does not teach a number of aspects of the claimed invention, including a supervising terminal connected

to the storage, and the storage logically partitions the plurality of interfaces, the plurality of disk drives, and the control unit on the basis of the information to be inputted to the supervising terminal.

US 5,659,786, George et al., discloses dynamic reconfiguration of system resources in a logically partitioned system without the need for operator involvement to free up resources. In operation, when started by an external stimulus, such as an operator command or a time-driven event, a hardware policy or PR/SM operator requests a physical configuration change. The processor controller element passes the request to LPAR, which translates the request into a request (or requests) to a logical partition (or partitions) to free up logical resources. LPAR sends the translated requests to operating systems in the logical partition(s), which respond as they would to an operator request by performing logical deconfiguration, and then physical deconfiguration via a signal to LPAR. LPAR, which may initiate deconfiguration requests to different partitions in parallel, evaluates the actions by each partition and, if necessary, consults a policy to make needed adjustments to insure that all needed resources are obtained. Finally, LPAR

sends the appropriate physical reconfiguration request(s) to the processor controller element for execution. (See, e.g., Column 2, lines 37-55.) Thus, George et al. do not teach the present invention, wherein a control unit in a storage system logically partitions a plurality of interfaces, a plurality of disk drives, and the control unit, and causes the partitioned plurality of interfaces, the partitioned plurality of disk drives, and the partitioned control unit to operate independently as a plurality of virtual storages.

US 5,704,055, George et al., discloses a data processing system that has a processing unit and a memory which provides a common pool of physical storage. This storage is initially assigned as either main storage or expanded storage during power-on. Subsequent to the initial assignment, storage assigned as main storage or expanded storage may be unassigned and thus returned to the common pool. Once returned to the common pool, the storage may be reassigned as either main storage or expanded storage. The storage reassignment is done dynamically without requiring a reset action and transparent to the operating system and any active application programs. (See, e.g., Abstract, and Column 5, lines 26-38.) Thus, George does not disclose the claimed invention, such as a

control unit that logically partitions a plurality of interfaces, a plurality of disk drives, and the control unit, and causes the partitioned plurality of interfaces, the partitioned plurality of disk drives, and the partitioned control unit to operate independently as a plurality of virtual storages.

US 5,790,852, Salm, discloses a computer system comprising a virtual storage, and means for organizing the virtual storage to provide storage space for parallel program execution in pre-allocated partitions. The virtual storage contains shared areas and private areas in a plurality of address spaces. Depending on the job to be performed, the organizing means allocates one or more dynamic partitions in the private areas in addition to the pre-allocated partitions, and the dynamic partitions are de-allocated after job termination to free the storage area for other use. (See, e.g., Column 1, lines 39-49.) Thus, Salm is not directed to a storage system per se, and does not disclose a storage system including a control unit that logically partitions a plurality of interfaces, a plurality of disk drives, and the control unit, and causes the partitioned plurality of interfaces, the partitioned plurality of disk drives, and the partitioned

control unit to operate independently as a plurality of virtual storages.

US 6,484,245, and US 6,742,090, Sanada et al., disclose a system that includes a storage controller 40 constituted from a fiber channel control unit 41 which may be a protocol processor including a direct memory access (DMA) for controlling data transmission between it and the host computers 10, 20, 30, a microprocessor 42 for controlling all possible operations of the storage controller, a control memory 43 for storing therein micro-programs for control of the operation of the controller along with control data associated therewith, a cache control unit 44 for controlling writing and reading data to and from the cache, a disk cache 45 for temporarily buffering write data and read data to/from a disk drive(s), a device interface control unit 46 which may be a protocol processor including DMA for controlling data transfer between it and its associative disk drives, and a panel 47 for use in inputting device configuration information to the storage controller, and a disk array subsystem 50 operable under control of the storage controller 40. (See, e.g., Column 4, line 58 - Column 5, line 24 of the '245 patent.) However, Sanda fails to teach the functionality of

the present invention, in which a control unit logically partitions a plurality of interfaces, a plurality of disk drives, and the control unit, and causes the partitioned plurality of interfaces, the partitioned plurality of disk drives, and the partitioned control unit to operate independently as a plurality of virtual storages.

US 6,606,690, Padovano, discloses a system in which a first SAN (Storage Area Network) server is configured to be coupled to a plurality of storage devices in a SAN via a first data communication network. The first SAN server allocates a first portion of the plurality of storage devices in the SAN to be accessible to at least one first host coupled to a second data communication network. The first SAN server allocates a second portion of the plurality of storage devices in the SAN to a first NAS server. The first NAS server configures access to the second portion of the plurality of storage devices to at least one second host coupled to a separate data communication network. (See, e.g., Column 2, lines 40-62.) Thus, Padovano is directed to servers that control configuration of storage, unlike the present invention, in which a storage system includes a control unit that logically partitions a plurality of interfaces, a

plurality of disk drives, and the control unit, and causes the partitioned plurality of interfaces, the partitioned plurality of disk drives, and the partitioned control unit to operate independently as a plurality of virtual storages.

US 6,640,278, Nolan et al., discloses a system for managing storage resources in a storage network according to storage domains. The system includes a plurality of communication interfaces, adapted for connection via communication media to clients and storage systems and the storage network. A processing unit is coupled with the plurality of communication interfaces and includes logic to configure a set of storage locations from the one or more storage systems in the network as a storage domain for a set of at least one client from the one or more clients in the storage network. (See, e.g., Column 2, lines 19-41.) Thus, Nolan is not concerned with the same problem as the present invention, and Nolan fails to teach a storage system that logically partitions a plurality of elements to operate as a plurality of virtual storages.

US 6,725,352, Goodman et al., discloses a method to partition a data storage and retrieval system into one or more

logical libraries, where that data storage and retrieval system includes a library controller, at least one data storage drive and at least one control port. Using Goodman's method, a data storage and retrieval system can be partitioned into smaller virtual libraries called logical libraries. A logical library comprises a subset of the complete physical data storage and retrieval system, where that logical library includes at least one data storage drive and at least one control port. (See, e.g., Column 2, lines 15-49.) However, Goodman does not teach the present invention, in which a storage system includes a control unit that logically partitions a plurality of interfaces, a plurality of disk drives, and the control unit, and causes the partitioned plurality of interfaces, the partitioned plurality of disk drives, and the partitioned control unit to operate independently as a plurality of virtual storages.

US 6,738,854, and US 6,763,419, Hoese et al., disclose a system in which collective storage is provided by storage devices 60, 62, and 64 which can have blocks allocated by programming means within a storage router 56. Virtual local storage is accomplished without limiting the performance of workstations 58 because the storage access involves native low

level, block protocols and does not involve the overhead of high level protocols and file systems required by network servers. Additionally, a supervisor unit 86 comprises a microprocessor for controlling operation of storage router 56 and to handle mapping and security access for requests between a Fibre Channel 52 and a SCSI bus 54. (See, e.g., Column 4, lines 7-59 of the '854 patent.) Thus, Hoese does not teach a storage system that logically partitions a plurality of elements to operate as a plurality of virtual storages.

US 6,754,776, Conway et al., discloses a system and method of logically partitioning a cache memory between computer domains using an extended memory address is disclosed. The extended memory address includes an address space identifier for extending a conventional memory address with at least one bit uniquely identifying the address space of a domain from which a data request is made to the cache memory. (See, e.g., Column 2, lines 28-34.) Thus, Conway does not disclose the claimed invention, in which a storage system includes a control unit that logically partitions a plurality of elements to operate as a plurality of virtual storages.

US 20020007366, Fontijn, discloses a method of implicitly partitioning the storage space available on a storage medium, into a storage medium for storing user data, and a recording device for storing user data on a storage medium. Thus, Fotijn does not teach the present invention, which includes a storage system having a control unit that logically partitions a plurality of interfaces, a plurality of disk drives, and the control unit, and causes the partitioned plurality of interfaces, the partitioned plurality of disk drives, and the partitioned control unit to operate independently as a plurality of virtual storages.

US 20020133539, Monday, discloses a system and method for attaching a remote storage device to a network. The remote storage device is detected and automatically incorporated into a new or existing logical storage volume. A size of a logical storage space provided by a file system is automatically expanded to include the additional storage capacity. The storage capacity is incorporated according to one or more pre-defined policies set by a system administrator or other user. To incorporate the storage capacity, a network-based connection is formed with the remote storage device such that a logical volume manger can configure the storage device as if

it were locally attached. For example, a physical volume is created for the remote storage device and the physical volume is added to a volume group based on the pre-defined policies. (See, e.g., paragraphs [0008]-[0009].) Thus, while Monday teaches the automatic reconfiguration of logical storage upon the addition of a new storage device, Monday does not teach the present invention, which includes a storage system having a control unit that logically partitions a plurality of interfaces, a plurality of disk drives, and the control unit, and causes the partitioned plurality of interfaces, the partitioned plurality of disk drives, and the partitioned control unit to operate independently as a plurality of virtual storages.

US 20030065898, Flamma et al., discloses a system for managing the storage of objects in a storage system having a plurality of different storage media divided into different partitions. The system includes a storage processor for determining whether a particular storage partition has reached a predetermined capacity threshold, and a data migration processor for identifying within the particular storage partition an object to be moved and for identifying a target destination partition for the particular object in response to

the capacity determination. The data migration processor identifies the target destination partition based on one or more of (a) media type of the particular storage partition, and (b) information identifying related objects in the target destination partition. The system also includes a transfer processor for transferring data representing the particular object to the target destination partition. (See, e.g., paragraph [0011].) Thus, while Flamma provides a system for migration of data within storage partitions, Flamma does not teach the present invention in which a storage system includes a control unit that logically partitions a plurality of elements to operate as a plurality of virtual storages.

US 20030120751, Husain et al., discloses system and method for managing data storage for a plurality of computer systems. The computer systems may be configured to access virtual network attached storage, but Husain fails to teach the present invention, and does not disclose a storage system that includes a control unit that logically partitions a plurality of elements to operate as a plurality of virtual storages.

US 20040111580,Weber et al., discloses systems and methods for managing requests of a host system to physical storage partitions. A storage system includes a plurality of storage elements with each storage element configured for providing data storage. A communications switch is communicatively connected to the storage elements for transferring requests to the physical storage partitions. A host system includes a storage router for mapping a portion of the physical storage partitions to logical storage partitions such that the host system can directly access the portion via the requests. Each of the storage elements includes a storage controller configured for processing the requests of the host system. The storage elements also include any of a disk storage device, tape storage device, CD storage device, and a computer memory storage device. (See, e.g., paragraphs [0010] and [0012].) However, Weber does not teach a storage system having a control unit that logically partitions a plurality of interfaces, a plurality of disk drives, and the control unit, and causes the partitioned plurality of interfaces, the partitioned plurality of disk drives, and the partitioned control unit to operate independently as a plurality of virtual storages.

US20030097393 and JP2003-157177, Kawamoto et al.,

disclose a virtual computer system having a hypervisor (allocating means) that divides a physical computer into a plurality of logical partitions (LPARs), that runs an operating system (OS) in each LPAR, and that controls allocation of resources of the physical computer to the LPARs. The virtual computer system consists mainly of a user interface, a load measuring means, and an adaptive control means. The user interface enables entry of one setting or a plurality of settings concerning the control actions of the virtual computer system. The load measuring means measures loads to be accomplished by the OSs in the LPARs. The adaptive control means (allocation ratio varying means) determines the allocation ratios of the computer resources relative to the LPARs according to the settings entered through the user interface and the loads to be accomplished by the OSs in the LPARs which are measured by the load measuring means. If the determined allocation ratios are different from the previous ones, the adaptive control means instructs the hypervisor to vary the allocation ratios. Furthermore, the hypervisor includes a means for dynamically varying the allocation ratios of the computer resources relative to the LPARs in response to the instruction issued from the adaptive

control means. (See, e.g., paragraph [0013].) Consequently, Kawamoto teaches a virtual computer system capable of dynamically and optimally allocating computer resources to LPARs according to the loads to be accomplished by OSs that run in the LPARs and a knowledge of workloads running on the OSs. However, Kawamoto does not teach the present invention in which a storage system has a control unit that logically partitions a plurality of interfaces, a plurality of disk drives, and the control unit, and causes the partitioned plurality of interfaces, the partitioned plurality of disk drives, and the partitioned control unit to operate independently as a plurality of virtual storages.

CONCLUSION

The Applicants submit that the foregoing discussion demonstrates the patentability of the claimed invention over the closest-known prior art. Accordingly, the requirements of 37 CFR §1.102(d) having been satisfied, the Applicants request that this Petition to Make Special be granted and that the application be examined according to prescribed procedures set forth in MPEP §708.02 (VIII).

The Applicants prepared this Petition in order to satisfy the requirements of 37 C.F.R. §1.102(d) and MPEP §708.02 (VIII). The pre-examination search required by these sections was "directed to the invention as claimed in the application for which special status is requested." MPEP §708.02 (VIII). The search performed in support of this Petition is believed to be in full compliance with the requirements of MPEP §708.02 (VIII); however, the Applicants make no representation that the search covered every search area that may contain relevant prior art. It is always possible that prior art of greater relevance to the claims may exist. The Applicants urge the Examiner to conduct his or her own complete search of the prior art, and to thoroughly examine this application in view

of the prior art cited above and any other prior art that may be located by the Examiner's independent search.

Further, while the Applicants have identified certain portions of each cited reference in order to satisfy the requirement for a "detailed discussion of the references, which discussion points out, with the particularly required by 37 C.F.R. §1.111(b) and (c), how the claimed subject matter is patentable over the references" (MPEP §708.02(VIII)), the Examiner should not limit review of these documents to the identified portions, but rather is urged to review and consider the entirety of each reference.

Respectfully submitted,

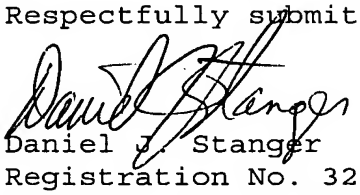

Daniel J. Stanger
Registration No. 32,846
Attorney for Applicants

MATTINGLY, STANGER & MALUR, P.C.
1800 Diagonal Rd., Suite 370
Alexandria, Virginia 22314
(703) 684-1120
Date: November 30, 2004